Developing a Regulatory Planning Information Retrieval Agent for Public Participation: Construction Method and Applications in Good-Governance Scenarios

Authors

WANG Shifu, LI Yonghao, HUANG Duo, LI Xinjian

Abstract

Against the dual background of the People-Centered City agenda and the deep digital transformation of territorial spatial planning, existing research has not yet adequately addressed the public need for precise, convenient, and low-cost intelligent retrieval of regulatory detailed planning information. This study first discusses a process-oriented pathway for cultivating planning agents, and then explains the construction logic of a planning information retrieval agent that integrates large language models (LLMs) and knowledge graphs within a retrieval-augmented generation (RAG) framework. On this basis, it proposes a three-stage method, using the Dify low-code development platform, for building ZoKnow, a public-participation-oriented regulatory planning retrieval agent. Seven statutory plans from Futian District, Shenzhen, including about 80,000 Chinese characters of planning text and 1,185 parcel-level control indicators, are used for application verification through prototype development, question-answering performance evaluation, and comparative analysis with the public version of the "One Map" platform. The results show that the proposed method is feasible and delivers superior retrieval performance. It is especially advantageous for terminology explanation, statutory-result traceability, and data analysis in public participation in regulatory planning. The study further discusses the potential of extending planning information retrieval agents toward multiple types of planning agents and their long-term significance for industry transformation. It provides a low-cost, reproducible, and scalable construction paradigm for planning agents that can support digital-intelligent planning transformation and technology-enabled spatial good governance.

Keywords

large language models; regulatory detailed planning; agent; knowledge graph; public participation; governance

Introduction

The People-Centered City concept emphasizes that cities are built by the people and for the people. The establishment of China's territorial spatial planning system marked a shift toward integrated, whole-domain, and all-factor development and protection [1-2]. Within this institutional transformation, regulatory detailed planning, hereafter referred to as regulatory planning, plays a pivotal role in linking higher-level territorial spatial planning with medium- and micro-scale development control. By defining land-use functions, development intensity, public service facility provision, and urban design control indicators, regulatory planning translates the goals of People-Centered City building into the material form and spatial distribution of benefits in specific places [3-4].

Ordinary citizens have rights to information, participation, and supervision throughout territorial spatial planning, construction, and governance. High-quality public participation is therefore essential to the effectiveness and fairness of regulatory planning preparation and implementation. Accurate access to regulatory planning content through the digitalization and intelligent processing of territorial spatial information is an important way to improve public participation [5-16]. Yet existing intelligent retrieval systems for regulatory planning remain insufficient in three respects: responding to public participation needs, improving functional convenience, and reducing data-processing costs [6,10-12,17-19]. First, next-generation digital platforms such as the territorial spatial planning "One Map" system and the China Spatial Planning Observation Network (CSPON) mainly serve government departments and planning professionals [6-9]. Second, public-facing digital platforms for regulatory planning retrieval still have limitations in rapidly obtaining query results, supporting the understanding of spatial meaning, and tracing statutory bases, which weakens the motivation and quality of public participation [10-12]. Third, regulatory planning outcomes, including planning texts and statutory atlases, are highly heterogeneous and unstructured; achieving accurate retrieval usually requires substantial data-processing input [20-22]. The development of LLMs, RAG, and knowledge graphs has brought new momentum to the digital transformation of planning technologies [23-25]. Early explorations in planning include master-plan knowledge graphs supported by code-based querying [26], knowledge-graph-based methods for dynamic urban cognition [24], and RAG-based planning text-generation assistants [27]. However, few studies have examined how to construct regulatory planning information retrieval agents that allow ordinary citizens to accurately obtain relevant information, quickly understand planning meanings, and operate at low cost. This gap constrains further improvement in public participation [26-28].

This study therefore proposes a construction method for ZoKnow, a public-participation-oriented regulatory planning information retrieval agent that integrates LLMs and knowledge graphs under a RAG framework, and verifies its application. It first discusses the process-oriented cultivation pathway of planning agents and explains the construction principles for regulatory planning retrieval agents. It then proposes a three-stage method for building ZoKnow. Using the texts and atlases of statutory plans in Futian District, Shenzhen as case data, the study develops a prototype, evaluates question-answering performance, and compares governance scenarios with the public version of the One Map platform. Finally, it identifies future directions for deepening the method and discusses the potential of planning information retrieval agents in good-governance applications and planning-industry transformation.

1 Cultivation Pathway for Planning Agents and Construction Principles for Regulatory Planning Information Retrieval Agents

1.1 Process-Oriented Cultivation Pathway for Planning Agents

An agent is a digital system with autonomy, reactivity, proactivity, and sociality. Based on multi-source data and advanced algorithms, it can perceive its environment, form memory, make decisions, interact with humans, and take actions to achieve defined objectives [29]. As LLMs have shown strong general capacity in natural-language understanding and generation, agents have evolved from fixed modular interaction systems into autonomous systems capable of multi-step action in complex environments. With LLMs as their reasoning and decision-making centers, agents can use planning, memory, and external tool calls to continuously acquire information, revise strategies, and complete tasks during interaction [30]. Current scholarship classifies LLM-based autonomous agents into general-purpose, tool-based, simulation, embodied, game, network, and assistant agents according to application purpose, functional boundary, and construction logic [30].

Planning agents are intelligent systems oriented to the full process of plan preparation, approval, amendment, implementation, and supervision. They can autonomously perceive, memorize, decide, interact, and execute. Combining the technical logic of LLM-based agents with practical experience in planning, a preliminary process-oriented cultivation pathway has emerged. This pathway includes scenario-demand identification, knowledge-base construction, data-base development, planning-tool adaptation, task-process orchestration, and feedback-based evaluation and governance. The knowledge base is the foundation for building a specific planning agent. It integrates professional knowledge and builds high-quality corpora so that planning knowledge can be structurally represented, semantically aligned, and made retrievable. The data base connects multi-source heterogeneous data, supports multidimensional perception and dynamic characterization of spatial elements, and strengthens the agent's understanding of complex urban conditions. Planning-tool adaptation embeds planning methods into the agent as callable digital function modules in the form of algorithms, protocols, or tools. Task-process orchestration enables coordinated perception, evaluation, analysis, and decision-making across scenarios. Finally, feedback and evaluation support continuous validation and optimization, improving reliability and effectiveness in practice.

1.2 Core Construction Principles of a Regulatory Planning Information Retrieval Agent

A regulatory planning information retrieval agent is an intelligent system that responds to public participation needs for regulatory planning information. It provides low-threshold interaction, task planning and orchestration, professional knowledge storage, external tool invocation, and executable planning tasks. Although ordinary citizens often have limited planning knowledge and technical capacity, they still need to efficiently obtain, understand, and analyze regulatory planning information for high-quality public participation. A tool-based planning agent is therefore an appropriate technical route. Its core feature is to call external tools to complete executable planning tasks. The key task can be summarized through a "3W" framework: whether to call a tool, which tool to call, and how to use the tool. This ensures that appropriate external tools are invoked at the right time for professional planning Q&A, data querying, and related tasks.

Drawing on tool-based agent construction logic [30] and information-system development requirements, the core architecture of a regulatory planning information retrieval agent can be divided into four components: the task and interaction system, or interaction layer; the planning and control system, or task layer; the knowledge and memory system, or knowledge layer; and the action and tool system, or engine layer (Figure 1).

The task and interaction system forms the input-output boundary through which the agent exchanges information with the external environment. Its core capability is to transform users' natural-language expressions into executable task descriptions. The planning and control system organizes decision-making and task control during execution. Under a preset reasoning framework, it decomposes complex problems into actionable steps, and uses evidence and feedback during action to trace back, correct, and supplement tasks. The knowledge and memory system combines the knowledge base and memory-management mechanisms needed for information acquisition. On the one hand, the agent must connect to updateable external knowledge sources to obtain reliable factual bases for specific tasks. On the other hand, it must maintain and read task states and key information across dialogue rounds so as to better understand execution status and users' actual needs. The action and tool system consists of functional modules for external tool invocation and execution. It externalizes key steps such as calculation and verification into executable tool chains, transforming analysis into a reviewable execution process.

Together, these four systems form an interdependent and dynamically feedback-based operating system. The task and interaction system clarifies "what problem is to be solved"; the planning and control system organizes "through what steps the problem should be solved"; the knowledge and memory system provides "on what evidence the problem should be solved"; and the action and tool system implements "by what means the problem should be solved." They enable a regulatory planning retrieval agent to move from simple interactivity toward complex task-solving capacities that are understandable, verifiable, and analyzable.

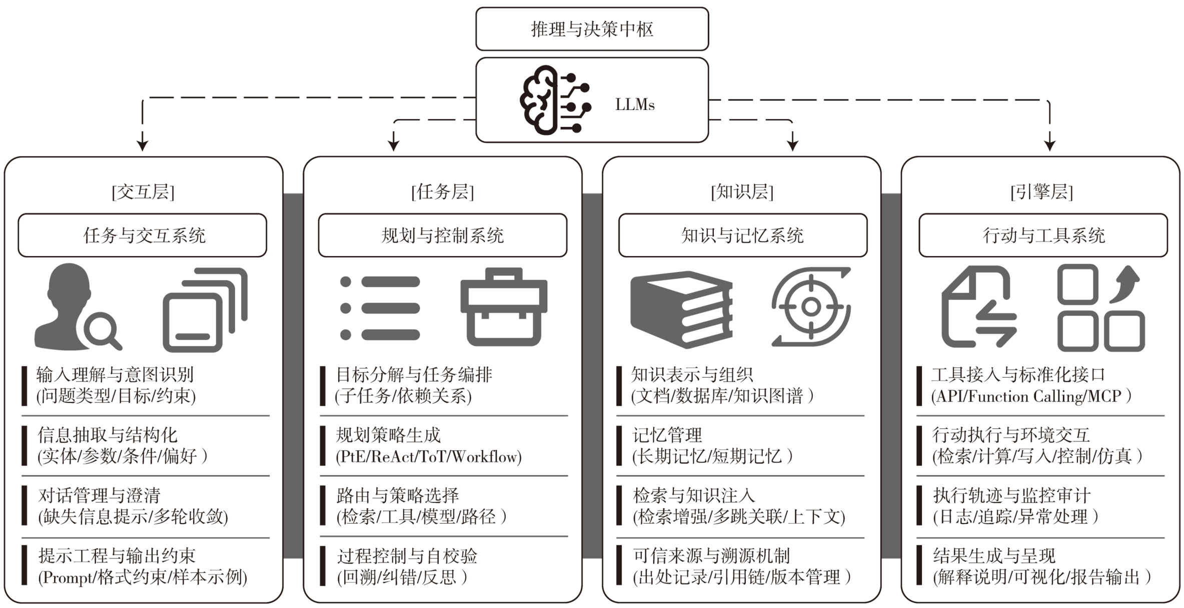

Figure 1. Four-layer core architecture for a regulatory detailed planning information retrieval agent.

1.3 Technical Logic for Embedding Knowledge Graphs into the Knowledge Layer under the RAG Framework

For a regulatory planning information retrieval agent, embedding a knowledge graph into the knowledge layer under a RAG framework offers three technical advantages: high accuracy, low operating cost, and convenient development. First, RAG introduces planning texts as external knowledge sources and retrieves them through semantic similarity, thereby providing LLMs with more sufficient and controllable contextual information. This can, to a certain extent, reduce hallucination in conceptual explanation and question-answering related to regulatory planning. Second, knowledge graphs store knowledge as "entity-relation-entity" triples and form structured databases. When embedded in the agent's knowledge layer, they can use graph retrieval and reasoning-chain construction to significantly improve the accuracy and controllability of LLMs in knowledge Q&A and multi-hop reasoning. Third, without pre-training, LLMs can automatically extract "entity-relation-entity" triples from planning outputs through prompts, accelerating knowledge-graph construction. They can also serve as efficient interfaces for graph-database querying and reasoning by converting natural language into query statements, thus enabling interactive knowledge acquisition between citizens and the agent.

2 Construction Method for a Public-Participation-Oriented Regulatory Planning Information Retrieval Agent

The regulatory planning information retrieval agent for public participation is a typical tool-based planning agent. Its function is to use clearly bounded regulatory information and data as external sources, provide ordinary citizens with accurate, fast, and low-threshold retrieval, checking, and explanation support, and transform outputs into planning evidence for deliberation. Based on this functional positioning, the construction principles above, and the Dify low-code development platform, this study proposes a three-stage method for building ZoKnow.

2.1 Preprocessing Regulatory Planning Texts and Atlases

For regulatory planning texts, Chinese is a paratactic language with relatively few explicit grammatical markers and strong contextual dependence in semantic relations. HanLP is therefore used to supplement natural-language syntactic information and reduce erroneous triple extraction caused by ambiguity in LLM processing. For regulatory planning atlases, which already have a certain degree of structure, image-cell segmentation is used under the condition that chart content remains complete and readable. This reduces the information density of each input and helps LLMs more clearly identify indicators, objects, and constraint relations in charts and tables.

2.2 Automated Knowledge-Graph Construction Based on LLMs

After multimodal data preprocessing, an LLM-centered workflow is constructed to automate the building of the regulatory planning knowledge graph. The knowledge-graph construction task is decomposed into named entity recognition, relation extraction, and label assignment. This allows unstructured regulatory planning text and chart information in statutory plans to be stably converted into structured information and then organized as a graph database.

2.3 Building ZoKnow on Dify with the Four-Layer Core Architecture

Following the construction principles of the regulatory planning retrieval agent, ZoKnow is designed on the Dify platform according to the core architecture of the interaction layer, task layer, knowledge layer, and engine layer. The regulatory planning knowledge graph is embedded into the knowledge layer to support precise information retrieval, evidence organization, and explainable answer generation (Figure 2).

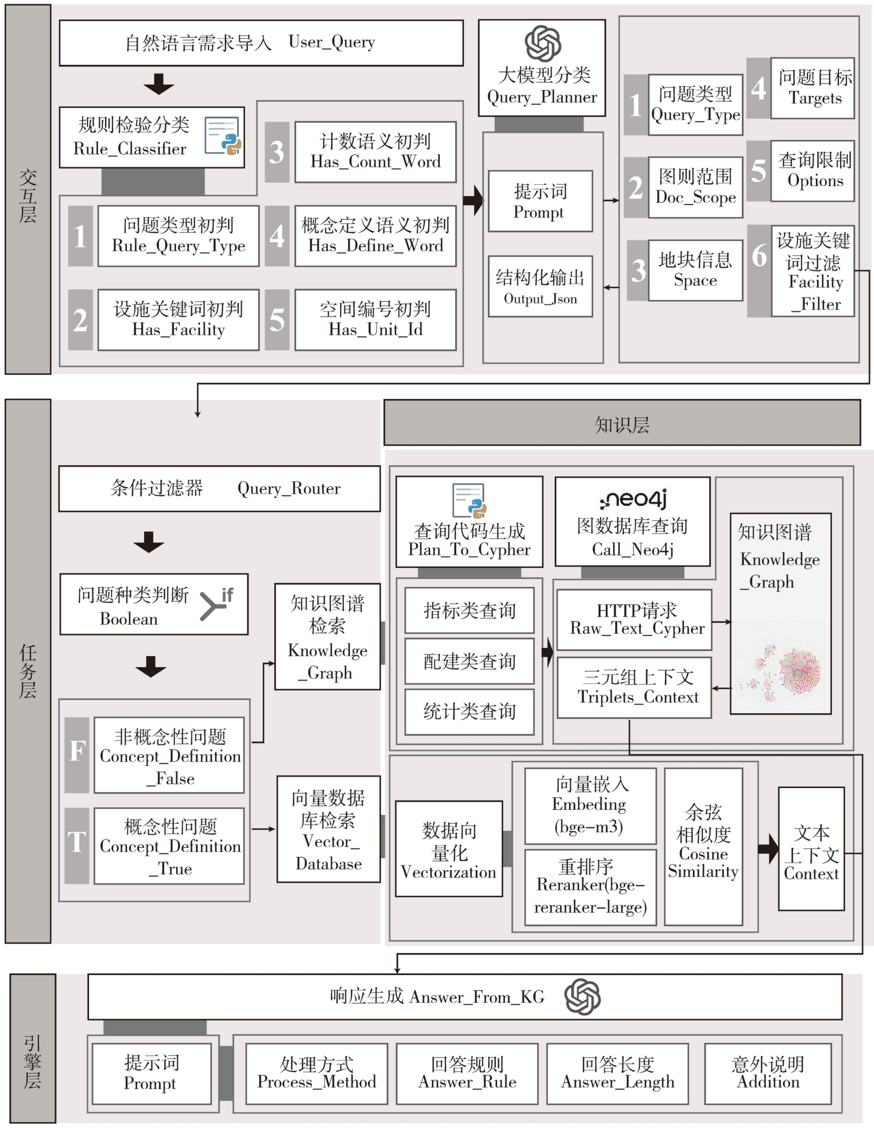

Figure 2. Four-layer framework of ZoKnow.

3 Application Verification Data and Methods

3.1 Case Data Selection and Sources

3.1.1 Case Data Selection

The application verification uses statutory plans for detailed planning urban units in Futian District, Shenzhen. After more than 20 years of continuous optimization, Futian District's statutory plans have developed a nationally advanced practical model in preparation rules, content structure, and expression formats, making them a representative case [18].

3.1.2 Data Sources and Types

The data are mainly drawn from statutory plans and related outcomes for Futian District publicly released by the Shenzhen Municipal Planning and Natural Resources Bureau. They include two main types.

First, statutory-plan atlases and tables are provided in PDF or image formats. They record parcel numbers, land-use codes, land-use functions, floor area ratio, building density, building height, green-space ratio, and public service facility provision requirements. These constitute the "hard-constraint" information source in the regulatory planning knowledge graph, expressed as mandatory control requirements [31].

Second, statutory-plan texts include explanation reports and amendment notes for each urban unit, in formats such as PDF and HTML. In addition to the "hard-constraint" information repeated in the atlases, these texts include planning background, planning objectives, and urban design control requirements. They therefore provide "soft-constraint" information sources in the regulatory planning knowledge graph, expressed as guidance requirements [31].

Because some statutory planning results were released long ago, their texts and atlases could not meet the data accuracy required for this study, and some content overlapped with later updated statutory plans. The study therefore manually screened the statutory plans of detailed planning urban units in Futian District and ultimately selected seven statutory plans from 2010 onward. The final dataset includes nearly 80,000 Chinese characters of planning text and control indicators for 1,185 parcels in 118 blocks contained in the planning atlases.

3.2 Application Verification Steps and Evaluation Methods

After the ZoKnow prototype was developed using the Futian District case data, application verification was conducted in two stages: comparative evaluation of retrieval performance and comparative analysis of governance scenarios (Figure 3). The first stage compares different retrieval strategies under a unified evaluation framework to verify ZoKnow's retrieval-performance advantage. The second stage examines its capacity to support public participation in real governance scenarios.

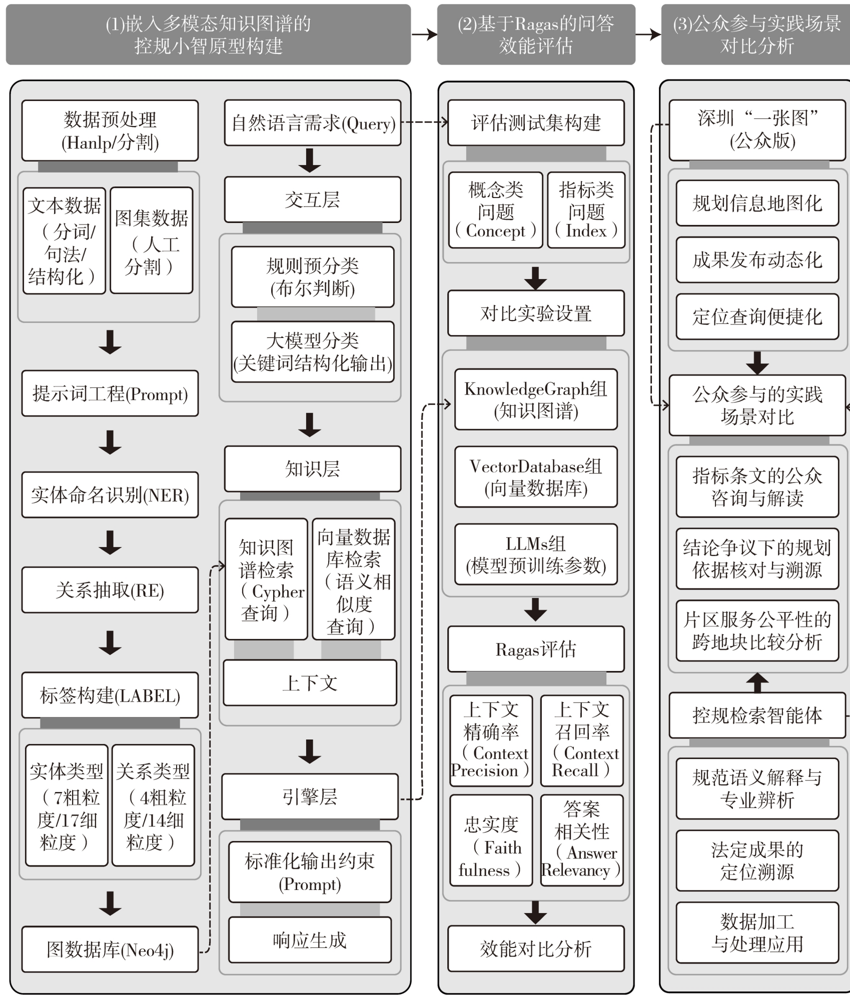

Figure 3. Steps of application verification.

3.2.1 Comparative Evaluation Method for Agent Retrieval Performance

To assess the practical performance of ZoKnow's core retrieval functions in real planning scenarios, this study constructed a structured Q&A evaluation test set based on regulatory planning information retrieval needs. For each of the seven statutory plans, 30 Q&A pairs were designed, forming 210 pairs in total for systematically testing the comprehensive performance of different retrieval schemes. In terms of content, the test set referred to current evaluation principles in regulatory planning, considered the core tasks of "qualitative determination, quantitative control, and spatial-form definition" and the preparation requirements of compliance, connectivity, and implementability, and simulated information retrieval needs that arise in real planning scenarios. In terms of rules, the construction of Q&A pairs strictly followed the six core rules of the Ragas (Retrieval-Augmented Generation Assessment System) evaluation framework, ensuring scientific structure and evaluation validity.

Based on the Ragas framework, the study systematically compared the performance of LLMs combined with different retrieval strategies in regulatory planning Q&A tasks. ZoKnow uses an LLM plus knowledge-graph scheme. Two control groups were established: a vector database scheme and a pure LLM scheme. The vector database scheme uses only the vector database of statutory plans as its external data source. The pure LLM scheme connects to no external data source and relies only on its pretrained parameters to generate answers. Four typical indicators were selected for evaluation: context precision, context recall, faithfulness, and answer relevancy.

3.2.2 Comparative Analysis Method for Governance Scenarios

This part focuses on three key needs for digital tools to improve the convenience of public participation: easier understanding, more accurate traceability, and more comprehensive analysis. It compares the performance of the public version of Shenzhen's regulatory detailed planning One Map platform with ZoKnow in three typical governance scenarios: public consultation and interpretation of indicators and clauses, verification and traceability of planning bases under disputes, and cross-parcel comparative analysis of district-level public service equity. This multidimensional comparison evaluates ZoKnow's application capacity.

In the scenario of public consultation and interpretation of indicators and clauses, ordinary citizens often need to query the floor area ratio, height limits, land-use function, and related clauses of a parcel when choosing where to buy a home, communicating within a community, or paying attention to nearby development. They often face a situation in which information can be found but is difficult to understand. Plain-language explanation and conceptual clarification are therefore needed so that citizens can quickly form understandable reference judgments.

In the scenario of checking and tracing planning bases, ordinary citizens involved in facility-location disputes or construction-process supervision often encounter inconsistent claims from different channels and lack verifiable evidence for deliberation. Retrieval systems therefore need to connect conclusions to specific planning texts, indicator tables, or document sources, helping citizens efficiently locate and obtain the planning information they need and improving the quality and efficiency of public participation.

In the scenario of cross-parcel comparison of public services, citizens discussing whether public services are balanced or whether facilities are insufficient often need to aggregate, compare, and conduct simple statistics across multiple parcels or districts. They also need results that can generate meaningful public-participation issues. Cross-parcel indicator statistics, difference comparison, and threshold checking can thus transform scattered queries into basic analytical results for public deliberation.

4 Application Verification Results

4.1 Comparative Evaluation of ZoKnow Retrieval Performance

The comparative results of retrieval performance under the Ragas framework are shown in Table 1.

Table 1. Evaluation of response performance of retrieval schemes based on the Ragas framework.

Evaluation type	Retrieval scheme	Context precision	Context recall	Faithfulness	Answer relevancy
Overall performance	ZoKnow (knowledge graph scheme)	0.92	0.84	0.99	0.68
Overall performance	Vector database scheme	0.64	0.57	0.99	0.40

Overall performance	Pure LLM scheme	0.19	0.12	0.99	0.00
Indicator-type questions	ZoKnow (knowledge graph scheme)	0.97	0.92	0.98	0.70
Indicator-type questions	Vector database scheme	0.37	0.31	0.98	0.26
Indicator-type questions	Pure LLM scheme	0.01	0.01	0.99	0.00
Concept-type questions	ZoKnow (knowledge graph scheme)	0.88	0.76	0.99	0.65
Concept-type questions	Vector database scheme	0.91	0.82	0.99	0.54
Concept-type questions	Pure LLM scheme	0.37	0.24	0.99	0.00

For overall performance, ZoKnow shows clear advantages in retrieval precision and focused response. Except for faithfulness, ZoKnow outperforms both the vector database scheme and the pure LLM scheme in context precision, context recall, and answer relevancy. This indicates that after the knowledge graph is introduced, ZoKnow can more accurately identify retrieval needs and more fully cover key clauses and indicator information required for answering. By contrast, the vector database mainly retrieves relevant fragments through semantic similarity; in complex scenarios where clauses and indicators are mixed, it is prone to redundant results and has difficulty expressing structured control indicators. The pure LLM scheme lacks support from external professional data and performs poorly across indicators. This result shows that, in domain-specific planning Q&A, connecting to external data sources is a key condition for improving retrieval accuracy.

For concept-type questions, vector retrieval has an advantage in semantic coverage, but ZoKnow remains more stable in focused response. Vector retrieval based on semantic similarity can cover part of the explanation of planning terminology and related clauses. However, ZoKnow still leads in answer relevancy. This reflects that, even when both schemes use vector data, ZoKnow supplements semantic retrieval with entity-relation information and structured constraints introduced by the knowledge graph. This helps the system organize answers around the core of the question and reduce overgeneralized explanations and irrelevant additions.

For indicator-type questions, ZoKnow has a particularly strong advantage and this supports its leading overall performance. The performance gap among schemes is most significant in this category. ZoKnow maintains obvious advantages in context precision, context recall, and answer relevancy, showing that the knowledge-graph scheme is more stable and explainable when processing highly structured information such as parcel numbers, control indicators, and numeric attributes. This is the main reason for its overall superiority. By contrast, the vector database scheme declines significantly in indicator-question evaluation, indicating that pure vector similarity struggles to reliably capture correspondence among graphical and tabular information. The result further confirms that, without structured external data support, LLMs are not well suited to high-accuracy and traceable regulatory planning indicator retrieval.

Overall, compared with the control schemes, ZoKnow demonstrates low-hallucination and high-accuracy characteristics in precise response, indicator retrieval, and focused answering. It is especially suitable for regulatory planning tasks centered on indicator verification and clause checking, and provides a technical foundation for supporting effective public participation.

4.2 Comparative Analysis of Governance Scenarios: ZoKnow and the Public Version of One Map

For public-participation governance needs such as consultation and interpretation of indicators and clauses, checking and tracing planning bases, and cross-parcel comparative analysis of district-level public services, the public version of Shenzhen's regulatory detailed planning One Map platform provides only basic display, search, and information listing functions (Figure 4). ZoKnow, by contrast, provides more complete functions for information interpretation, traceability, and analysis, and can effectively serve specific governance needs (Figure 5).

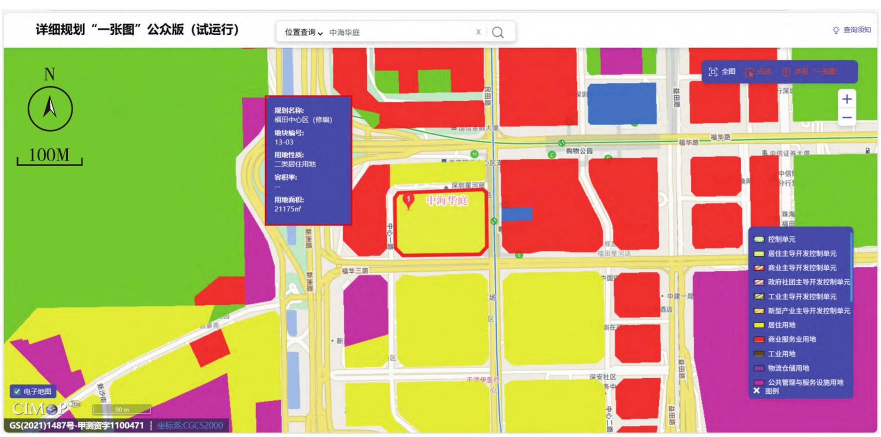

Figure 4. Public version of the One Map platform for Shenzhen's regulatory detailed plans.

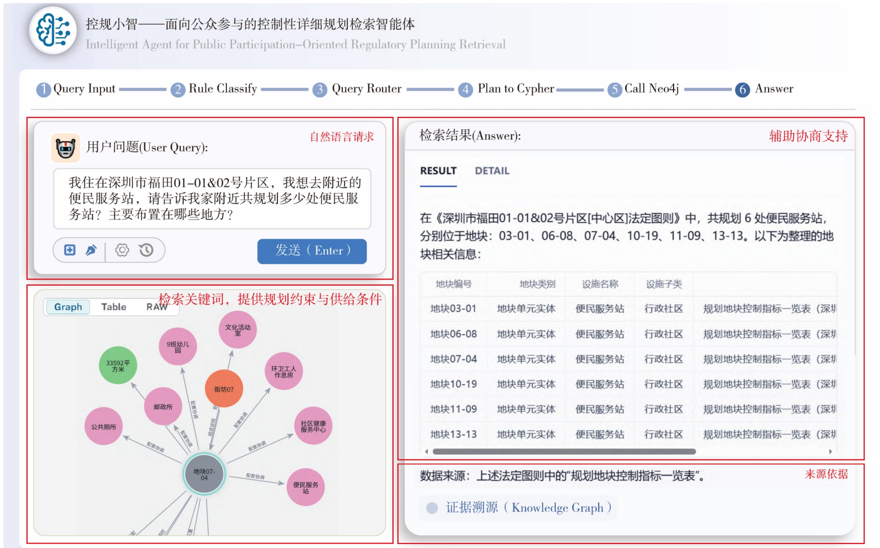

Figure 5. Selected governance applications of ZoKnow.

First, ZoKnow provides understandable interpretation of regulatory semantics and professional concepts. The threshold for public participation in regulatory planning does not only lie in whether parcel indicators or clauses can be retrieved, but also in whether highly specialized expressions in regulatory planning outcomes can be translated into semantic content that ordinary citizens can understand. Shenzhen's One Map public version can already retrieve floor area ratio, land-use function, building height limits, and related information, but it does not explain the planning meaning behind these indicators. ZoKnow explains professional terms and distinguishes concepts while returning retrieval results, producing answers that better match citizens' understanding habits and lowering the threshold for comprehension.

Second, ZoKnow provides verifiable statutory-result checking and traceability. In public participation scenarios such as facility layout, implementation supervision, and adjustment feedback, ordinary citizens often face problems such as "conclusions without evidence" or "evidence that is hard to locate." They may lack a verifiable evidence chain, or spend substantial time comparing their claims with relevant planning documents. Compared with the query and display functions of Shenzhen's One Map public version, ZoKnow simultaneously provides structured statutory-result provenance when outputting results, including corresponding text paragraphs, indicator tables, and document sources. This supports evidence-based deliberation and responsibility identification in public participation, enabling citizens to participate in negotiation with grounds and rules to rely on, and improving the efficiency and quality of participation.

Third, ZoKnow provides analyzable cross-district public-service data calculation. In regulatory planning participation, ordinary citizens need not only simple parcel information but also comparative analysis across parcels, such as whether public services are balanced within a district, whether facility supply has gaps, and whether indicators differ abnormally across parcels. The One Map public version has limited data-processing capacity and largely remains at the level of simple information listing. ZoKnow can use LLMs to translate natural-language needs, conduct structured retrieval and calculation, summarize indicators, compare differences, check thresholds, and generate statistical outputs. It therefore helps citizens better analyze and understand planning data and district development characteristics.

5 Conclusions and Outlook

5.1 Summary of the Construction Method and Directions for Deepening

This study discusses the process-oriented cultivation pathway of planning agents and explains the construction principles of a public-participation-oriented regulatory planning information retrieval agent. On this basis, and using the Dify platform, it proposes a three-stage method consisting of text and atlas data processing, automated knowledge-graph construction, and the building of a core architecture composed of interaction, task, knowledge, and engine layers. Application verification based on Futian District, Shenzhen shows that the regulatory planning retrieval agent constructed in this study significantly outperforms the traditional LLM-plus-vector-retrieval scheme and the pure LLM scheme in context precision, recall, and related indicators, especially in strongly constrained indicator Q&A. Compared with the public version of Shenzhen's regulatory detailed planning One Map platform, the agent enables citizens to use natural language to explain regulatory planning terminology, query parcel indicators, and trace statutory bases, thereby lowering the technical threshold for understanding and participating in regulatory planning affairs.

Future research should deepen the method in three respects. First, the agent's application scenarios should be expanded to verify its generalization capacity in complex situations involving concrete spatial interest conflicts. Second, autonomous iterative updating should be promoted so that the agent can adaptively evolve in response to real public participation needs and planning knowledge updates. Third, the coupling mechanism between information retrieval and deliberative consensus-building should be explored, further empowering multi-stakeholder participation. These directions will help the agent evolve from a single application for regulatory planning information retrieval into an open governance tool embedded in public participation, and may form a transferable and reusable construction paradigm for public urban knowledge infrastructure.

5.2 Technical Positioning of Planning Information Retrieval Agents and Expansion of Good-Governance Applications

The regulatory planning information retrieval agent proposed in this study is a specific branch of tool-based planning agents in regulatory detailed planning. Its core feature is to convert multi-source heterogeneous professional knowledge in planning into general data that the public can understand and machines can recognize. Similar planning information retrieval agents can serve as the knowledge base in the process-oriented cultivation pathway of planning agents, or be embedded as sub-agents in general-purpose, assistant, simulation, and other types of planning agents. They can provide high-precision, traceable, and iteratively updateable professional knowledge, thereby strengthening decision support in planning practice.

Beyond public-participation-oriented regulatory planning information retrieval, planning agents are gradually becoming important technical carriers that connect professional knowledge, multi-source data, digital tools, and different service users. They can support governance scenarios such as outcome review, preparation assistance, spatial analysis, facility allocation, scheme simulation, urban sensing, participatory deliberation, and planning education. Relevant cases show that the application boundary of planning agents is expanding from single information retrieval and Q&A toward more complex practical processes, including tool invocation, task orchestration, scheme generation, process simulation, and decision support, thereby supporting the realization of good-governance objectives.

5.3 Long-Term Significance of Planning Agents for Industry Transformation under the Idea of Good Governance

Urban good governance emphasizes plural inclusion, joint consultation, co-construction, and shared benefits. From this perspective, the development and application of planning agents can push traditional planning toward open social practice. They not only improve the precision and efficiency of planning information processing and decision judgment, but also reshape the generation, interaction, and use of planning knowledge. They are not simply the "digital twins" of planners, but inclusive governance-support platforms for multi-stakeholder participation, enabling digitally and intelligently empowered pathways for professional empowerment, data support, institutional adaptation, and consensus formation.

Planning agents will profoundly transform planning workflows and planners' competency systems. In the future, agents may be embedded throughout plan preparation, approval, amendment, implementation, and supervision, substantially improving the automation and intelligence of planning work, enhancing scientific rigor and implementability, and saving labor costs. Planners' core capabilities must also evolve. Building on traditional scheme design, plan preparation, and management coordination, planners need stronger capabilities in learning and applying AI tools and data processing, and should integrate these tools into the whole process through which planning serves urban good governance, thereby reshaping professional value.

Empowering the People-Centered City and good governance through planning agents will be a long-term undertaking for the planning industry. On the one hand, AI is developing rapidly, and models, algorithms, and platforms continue to update, greatly accelerating iteration among technologies. On the other hand, the reconstruction of relationships between people and technology, and among people themselves, remains gradual and is constrained by social ethics, governance structures, public cognition, and other factors. Planning decisions still require dynamic balancing among social equity, public interest, and the protection of private rights. The deep tension between technology and society means that planning agents are not merely technical tools, but a form of future-oriented spatial good-governance infrastructure. They can lower the threshold for information access, understanding, and analysis; reshape communication among government, the public, and the planning profession; and improve the inclusiveness and effectiveness of planning work. The key to fully using planning agents lies not only in the sophistication of technology, but also in whether effective connections can be built between rapidly evolving technical systems and relatively stable social institutions, enabling everyone in the city to become part of a place-based community that learns, influences, and progresses together.

Notes

[1] Hallucination in LLMs refers to generated content that is semantically plausible but inaccurate in terms of facts or regulatory clauses.

[2] The six rules for constructing Ragas Q&A pairs are: the answer can be fully derived from the given context; the information is not a simple fact; the answer contains no links; the difficulty is moderate; the question is reasonable and answerable by humans; and expressions such as "provide context" are not included.

[3] The public version of Shenzhen's regulatory detailed planning One Map platform publishes detailed planning outcomes through an electronic map. Its features include, but are not limited to, map-based

planning information, dynamic publication of results, and convenient location-based querying (https://pnr.sz.gov.cn/d-xgmap/).

[4] In the Ragas evaluation, the prompt required the model to avoid fabrication when retrieval evidence was absent, and Temperature was set to 0. The answer relevancy score of 0 in the pure LLM scheme means that the answer cannot be used to infer the original question; it is not an evaluation of factual correctness. In addition, because this study incorporated the question itself into the context, the high faithfulness scores in the three experiments only indicate that the answers were supported by the given context; they do not mean that the answers were necessarily factually correct.

References

[1] LIU Y S, ZHOU Y. Territory spatial planning and national governance system in China[J]. Land Use Policy, 2021, 102: 105288.

[2] YU Y D, CAO X D, YE D, et al. Critical considerations in provincial territorial spatial planning in China: a qualitative and spatial analysis[J]. Humanities & Social Sciences Communications, 2025, 12(1): 1422.

[3] WANG Yan, ZENG Xiangkun. Exploring detailed planning evaluation within urban development boundaries from a full-cycle management perspective[J]. Urban Planning Forum, 2024(5): 113-118.

[4] JIANG Tao, LI Yanxin, JIANG Mei. Research on the urban design control-element system at the regulatory detailed planning stage[J]. Urban Planning Forum, 2017(4): 65-73.

[5] LIU Xiaochang, WU Zhiqiang. Constructing a theoretical paradigm for AI-assisted territorial spatial diagnosis and transforming methods of thinking and practice[J]. Urban Planning Forum, 2026(1): 12-20.

[6] JIANG Huaxiong. Implementation dilemmas, solutions, and implications of Western planning support systems[J]. Urban Planning International, 2023, 38(1): 117-123.

[7] DARGAN S, KUMAR M, AYYAGARI M R, et al. A survey of deep learning and its applications: a new paradigm to machine learning[J]. Archives of Computational Methods in Engineering, 2020, 27(4): 1071-1092.

[8] JIANG H, LI M, WITTE P, et al. Urban chatter: exploring the potential of ChatGPT-like and generative AI in enhancing planning support[J]. Cities, 2025, 158: 105701.

[9] BATTY M. The digital transformation of planning[J]. Environment and Planning B: Urban Analytics and City Science, 2021, 48(4): 593-597.

[10] WILSON A, TEWDWR-JONES M, COMBER R. Urban planning, public participation and digital technology: app development as a method of generating citizen involvement in local planning processes[J]. Environment and Planning B: Urban Analytics and City Science, 2019, 46(2): 286-302.

[11] SHAHAB S, BAGHERI B, POTTS R. Barriers to employing e-participation in the Iranian planning system[J]. Cities, 2021, 116: 103281.

[12] ARANDA N R, DE WAEGEMAEKER J, VAN DE WEGHE N. The evolution of public participation GIS (PPGIS) barriers in spatial planning practice[J]. Applied Geography, 2023, 155: 102999.

[13] YANG Junyan, ZHANG Xun, SHI Yi, et al. Seeing objects and people: urban design exploration based on digital portraits of urban populations[J]. Urban Planning Forum, 2025(4): 41-49.

[14] CHEN Jun, ZHANG Shanqi, ZHANG Bing, et al. Development directions for empowering territorial spatial planning with spatiotemporal information[J]. Urban Planning Forum, 2025(2): 20-27.

[15] HERZOG O, PAN Haixiao, DENG Zhituan, et al. Empowering urban planning with next-generation artificial intelligence: opportunities and challenges[J]. Urban Planning Forum, 2023(4): 1-11.

[16] CHEN Yourong, ZHANG Lina. Reflections on the digital transformation of detailed planning under the territorial spatial planning system: a case study of the Nanhu area in Fuzhou[J]. Urban Planning Forum, 2025(4): 96-102.
[17] MA Chenxi, YANG Jianqiang, XU Zun, et al. Reform and adaptive adjustment of regulatory detailed planning in the context of territorial spatial planning[J]. Urban Development Studies, 2025, 31(1): 26-32.
[18] WU Songtao, WANG Jingyuan, WEN Shipeng, et al. Hierarchy, response, and coordination: evolution and reflection on detailed planning units oriented to meso-level tool optimization[J]. City Planning Review, 2025, 49(1): 97-107.
[19] WU Zhiqiang, SHEN Qingji, WANG Yajuan, et al. A new planning paradigm: concept, methodological system, and practical path[J]. Urban Planning Forum, 2025(6): 11-17.
[20] KACZMAREK I, IWANIAK A, SWIETLICKA A, et al. A machine learning approach for integration of spatial development plans based on natural language processing[J]. Sustainable Cities and Society, 2022, 76.
[21] DENG Z Y, FU X Y, LI C S. Automating urban policy extraction: a large language model-based framework for extracting local heat policies from planning documents[J]. Environment and Planning B: Urban Analytics and City Science, 2025.
[22] SILVENNOINEN H, CHADZYNSKI A, FARAZI F, et al. A semantic web approach to land use regulations in urban planning: the ontozoning ontology of zones, land uses and programmes for Singapore[J]. Journal of Urban Management, 2023, 12(2): 151-167.
[23] NIU Xinyi, LIU Sihan, SANG Tian, et al. Professional learning of large models: constructing an image-generation model incorporating knowledge of urban spatial morphology design[J]. Urban Planning Forum, 2025(1): 55-63.
[24] YAN Fengying, LU Yunmo. Urban dynamic cognition method and micro-scenario applications based on knowledge graphs[J]. Urban Planning Forum, 2025(5): 39-46.
[25] YANG Junyan, ZHU Xiao, SHAO Dian. Looking back at history: the century-long evolution and future trends of urban design technology based on knowledge graphs[J]. Urban Planning Forum, 2021(6): 20-27.
[26] ZHANG W S Q, CHEN Y M, LIU X P, et al. A large language model-based approach to building the knowledge graph for master plan: a case study in Guangzhou, China[J]. Land Use Policy, 2025, 159: 106395.
[27] ZHU H, ZHANG W J, HUANG N X, et al. PlanGPT: enhancing urban planning with tailored language model and efficient retrieval[J/OL]. ArXiv Preprint, 2024. https://arxiv.org/abs/2402.19273.
[28] WU Zhiqiang, YAN Juan, XU Haowen, et al. Ten annual key issues in the development of urban and rural planning, 2024-2025[J]. Urban Planning Forum, 2024(6): 8-11.
[29] WOOLDRIDGE M, JENNINGS N R. Intelligent agents: theory and practice[J]. Knowledge Engineering Review, 1995, 10(2): 115-152.
[30] WANG L, MA C, FENG X Y, et al. A survey on large language model based autonomous agents[J]. Frontiers of Computer Science, 2024, 18(6): 186345.
[31] ZHAO Guangying, LIU Shujuan. Speculative discussion on the logic of planning guidance and control from the perspective of administrative acts: compulsory content and rigid control[J]. Urban Planning Forum, 2025(5): 31-38.

AI-assisted translation